



AF/ 3625
PATENT

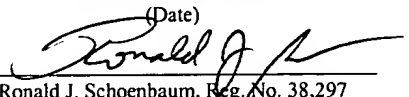
Case Docket No. AMAZON.012A1
Date: November 6, 2003
Page 1

In re application of : Bezos, et. al.
Appl. No. : 09/377,447
Filed : August 19, 1999
For : Computer Services For
Assisting Users in Locating
and Evaluating Items in an
Electronic Catalog Based
on Actions Performed by
Members of Specific User
Communities
Examiner : James Zurita
Art Unit : 3625

I hereby certify that this correspondence and all marked attachments are being deposited with the United States Postal Service as first-class mail in an envelope addressed to: Mail Stop INTERFERENCE, Board of Patent Appeals and Interferences, U.S. Patent and Trademark Office, P.O. Box 1450, Alexandria, VA 22313-1450, on

November 6, 2003

(Date)


Ronald J. Schoenbaum, Reg. No. 38,297

RECEIVED
DEC 02 2003
GROUP 3600

**MAIL STOP INTERFERENCE
BOARD OF PATENT APPEALS AND INTERFERENCES
UNITED STATES PATENT AND TRADEMARK OFFICE
P.O. Box 1450
Alexandria, VA 22313-1450**

Sir:

Transmitted herewith in triplicate is an Appellants' Brief to the Board of Patent Appeal, which includes an article attached as Appendix B. Also enclosed are:

- (X) A check in the amount of \$440.00 for the filing brief and the extension fee; and
- (X) A return prepaid postcard.

11/20/2003 CCHAU1 00000150 09377447

01 FC:1402 330.00 OP

11/20/2003 CCHAU1 00000150 09377447

01 FC:1251 110.00 OP

RECEIVED
NOV 10 PM 2:50
BOARD OF PATENT APPEALS
AND INTERFERENCES

PATENT

Case Docket No. AMAZON.012A1

Date: November 6, 2003

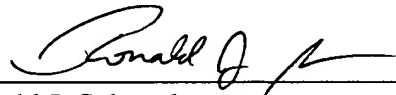
Page 2

An extension of time to respond for 1 month is hereby requested.

Time Extension Fee:

- ☒ (X) one month (\$110)
- ☐ () two months (\$420)
- ☐ () three months (\$950)

If applicant has not requested a sufficient extension of time and/or has not paid any other fee in a sufficient amount to prevent the abandonment of this application, please consider this as a Request for an Extension for the required time period and/or authorization to charge our Deposit Account No. 11-1410 for any fee which may be due. Please credit any overpayment to Deposit Account No. 11-1410.



Ronald J. Schoenbaum
Registration No. 38,297
Attorney of Record
Customer No. 20,995
(949) 760-0404

RECEIVED

NOV 12 2003

GROUP 3600

AMAZON.012A1

PATENT

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Applicants : Bezos, et al.) Group Art Unit: 3625
Appl. No. : 09/377,447)
Filed : August 19, 1999)
For : Computer Services for Assisting)
Users in Locating and)
Evaluating Items in an)
Electronic Catalog based on)
Actions Performed by Members)
of Specific User Communities)

Examiner : James Zurita

APPELLANTS' BRIEF
PURSUANT TO 37 C.F.R. § 1.192

Board of Patent Appeals and Interferences

Washington, D.C. 20231

Dear Sir:

Appellants, Applicants in the above-captioned patent application, appeal the final rejection of Claims 1-31 and 43-54, 56 and 57 set forth in the final Office Action mailed on July 9, 2003. A check for the filing fee is enclosed. Please charge any additional fees that may be required now or in the future to Deposit Account No. 11-1410.

I. REAL PARTY IN INTEREST

The real party of interest in the present application is Amazon.com, Inc.

II. RELATED APPEALS AND INTERFERENCES

No related appeals or interferences are pending.

Appl. No. : 09/377,447
Filed : August 19, 1999

III. STATUS OF CLAIMS

Claims 1-54, 56 and 57 are currently pending in the application, and are attached hereto as an appendix. All of the pending claims were finally rejected by the Examiner (note that Claim 55 was apparently rejected inadvertently, as it was canceled by Applicants' Amendment dated April 15, 2003).

Although Appellants believe all of the claim rejections are improper, Appellants are omitting Claims 32-42 from this appeal in order to reduce the number of issues on appeal. Appellants are also grouping the claims as set forth below in order to further reduce the issues on appeal. Appellants reserve the right to pursue the non-appealed claims in a continuing application.

IV. STATUS OF AMENDMENTS

Appellants did not amend any claims in response to the final Office Action.

V. SUMMARY OF THE INVENTION

The present application discloses various methods for providing users of an electronic catalog of items with information that is helpful to making informed item selection decisions. This information preferably includes data regarding associations between specific catalog items and specific user groups or "communities" (companies, clubs, cities, user groups, etc.), as detected by analyzing user activity data (e.g., purchase histories). For example, in one embodiment, when a particular item is determined to have a substantially higher popularity level within a particular user community than in a general user population, the relationship between this item and the user community is called to users' attention. This may be accomplished, for instance, by generating and displaying community-based popular items lists, and/or by sending email-based notifications to users (see Figures 2 and 4). Users of the electronic catalog can thus readily identify those items that have gained acceptance within, or that tend to distinguish, recognizable user communities. See present application at page 2, line 14 to page 3, line 22.

The user communities may include "explicit membership" communities that are explicitly "joined" by users of the system. For example, users of the system may be permitted to indicate their affiliations with specific colleges, local community groups, or professional organizations

Appl. No. : 09/377,447
Filed : August 19, 1999

(see Figure 1). The communities may additionally or alternatively include "implicit membership" communities that are not explicitly joined by users, but rather are formed based on information that is known about users. One type of implicit membership community that is disclosed is an email-based community that is formed based on user email addresses. For example, all users having email addresses ending in "microsoft.com" may be treated as belonging to the community "Microsoft.com Users." Implicit membership communities associated with specific geographic regions (cities, states, etc.) may be formed based on shipping addresses of users. See present application at page 7, lines 4-13.

To identify the catalog items that have attained disparate popularity levels within a particular user community, user activity data (e.g., purchase histories) of members of the community may be compared to the user activity data of a more general user population (e.g., all users of the system). One algorithm that may be used for this purpose is the censored chi-square algorithm detailed in the appendix of the present application; other algorithms that may be used are described elsewhere in the application. The identified items may be referred to as "characterizing" items, as they tend to distinguish members of the community from other users. See present application at page 8, line 28 to page 9, line 17, and pages 25-28.

Once one or more characterizing items have been identified for a given community, the relationship between these items and the community is preferably called to the attention of users of the electronic catalog. For example, users may be notified that a particular book about the Java programming language has attained a high or disparate popularity level among Microsoft Corporation employees. This information may be extremely helpful to a user who is seeking to purchase a book about this programming language. Users may be notified of the relationship via web pages of the electronic catalog (see Figure 2), email communications (see Figure 4), or other communications methods.

A commercial implementation of the invention known as *Purchase Circles* is accessible at the following URL: www.amazon.com/exec/obidos/subst/community/community.html/. The Purchase Circles service has inspired others in the computer science field to develop enhancements that facilitate the viewing and analysis of geographic-based Purchase Circles data. See DeMumbrum et al, *Visualizing Amazon.com's Purchase Circles*, College of Computing, Georgia Institute of Technology, attached hereto as Appendix B.

Appl. No. : 09/377,447
Filed : August 19, 1999

The "Contact Information Exchange" and "Purchase Notification" services disclosed in the present application are not pertinent to this appeal, and therefore are not discussed herein. Claims for aspects of these services are being pursued in U.S. Appl. No. 09/377,322, which was filed concurrently with the present application.

VI. ISSUES PRESENTED ON APPEAL

The following issue is presented:

Whether Claims 1-31 and 43-54, 56 and 57 are properly rejected under 35 U.S.C. § 103 as unpatentable over Anderson (U.S. Patent No. 5,974,396) in view of Bieganski (U.S. Patent No. 6,321,221), and in further view of Robinson (U.S. Patent No. 5,790,426).

VII. GROUPING OF CLAIMS

All of the rejected claims in the present application should not stand or fall together. Appellants, for purposes of this appeal only, are grouping the claims as follows:

GROUP 1 consists of independent Claims 1 and 12 and their respective dependent claims (collectively Claims 1-25). These claims are directed generally to a system and method for assisting users in selecting items from an electronic catalog. The independent system claim of this group refers to a computer process that (1) analyzes data to identify catalog items that have substantially higher popularity levels within particular communities relative to their respective popularity levels among a general user population, and (2) notifies users of the items and associated communities for which such popularity disparities exist to assist such users in selecting items from the electronic catalog.

GROUP 2 consists of independent Claim 26 and its dependent claims (collectively Claims 26-31). These claims are directed to an embodiment in which user email addresses are used to identify a subset of users, and in which at least purchase history data is processed to identify at least one item that characterizes this subset of users. Users are electronically notified of a relationship between the item and the subset of users.

GROUP 3 consists of independent Claim 51 and its dependent claims, except nos. 52 and 53 (collectively Claims 51 and 54-57). Independent Claim 51 is directed to an embodiment in which user email addresses are used to identify a plurality of organization-specific groups of

Appl. No. : 09/377,447
Filed : August 19, 1999

users, and in which user activity data is analyzed to generate organization-specific popular items lists. These organization-specific popular items lists are exposed to users of the electronic catalog to allow users to identify items that are popular within specific organizations.

GROUP 4 consists of dependent Claim 52, which further defines Claim 51 (Group 3) by reciting the following: “wherein exposing the organization-specific popular items lists to the users comprises providing to a target user a personalized web page that displays organization-specific popular items lists of organizations with which the target user is affiliated.”

GROUP 5 consists of dependent Claim 53, which further defines Claim 51 (Group 3) by reciting the following: “wherein analyzing the activity history data comprises substantially excluding from consideration activity history data falling outside a selected look-back period.”

GROUP 6 consists of independent Claim 43 and its dependent claims (collectively Claims 43-50). Claim 43 is directed to an embodiment in which user shipping addresses are used to identify the subset of users.

Appellants believe many of the dependent claims further define over the applied references; however, in the interest of administrative economy and efficiency, Appellants wish to narrow the issues by grouping the claims as set forth above.

VIII. ARGUMENT

A. EACH CLAIM GROUP IS PATENTABLY SEPARATE

For the reasons set forth below, each claim group is patentably separate.

Group 1: Both independent claims of Group 1 include limitations that provide patentable distinctions over the applied references. For example, both claims recite limitations related to the identification of items that have substantially higher popularity levels within specific communities; and include limitations regarding notifying users of a relationship between these items and communities to assist users in selecting items from the electronic catalog.

Group 2: Unlike the claims of Group 1, all of the claims of Group 2 require the identification of a subset of users that have email addresses that satisfy a particular criteria, and require the identification of at least one item that characterizes this subset of users. These limitations (within the context of the other claim limitations) are not suggested by the applied

Appl. No. : 09/377,447
Filed : August 19, 1999

references, and thus provide an additional patentable distinction over the applied references. Group 2 is therefore patentably separate from Group 1.

Group 3: The sole independent claim in Group 3 recites, for each of a plurality of organizations, identifying a respective group of users [of an electronic catalog] who are deemed to be affiliated with the respective organization by virtue of a domain name of the respective organization occurring within an email address of each such user, to thereby identify a plurality of organization-specific groups of users. The claim further recites, for each organization-specific group of users, analyzing activity history data to generate a list of items within the electronic catalog that distinguish the organization-specific group from a general population of the users, to thereby generate a plurality of organization-specific popular items lists. Further, the claim recites “exposing the organization-specific popular items lists to the users of the electronic catalog in association with the organizations to which such popular items lists correspond, to allow users to identify items that are popular within specific organizations.” The foregoing limitations are not included in Groups 1 and 2, and are not disclosed or suggested by the applied references. Group 3 is therefore patentably separate from Groups 1 and 2.

Group 4: The sole claim in Group 4 recites “providing to a target user a personalized web page that displays organization-specific popular items lists of organizations with which the target user is affiliated.” This feature is not required by the claims in Groups 1-3, and provides a separate patentable distinction over the applied references. Group 4 is therefore patentably separate from Groups 1-3.

Group 5: The sole claim in Group 5 recites “substantially excluding from consideration activity history data falling outside a selected look-back period.” This feature is not required by the claims in Groups 1-4, and provides a separate patentable distinction over the applied references. Group 5 is therefore patentably separate from Groups 1-4.

Group 6: Unlike the claims in Groups 1-5, the sole independent claim in Group 6 recites the use of user shipping addresses to identify a subset of users that reside in a common geographic region, and recites the processing of user activity data “to identify an item whose popularity level among the subset of users is substantially higher than a popularity level of the item among the population of users.” These limitations (in the context of the other claim language) are not

Appl. No. : 09/377,447
Filed : August 19, 1999

disclosed or suggested by the applied references, and provide a separate basis for patentability over the applied references. Group 6 is therefore patentably separate from Groups 1-5.

B. DISCUSSION OF REFERENCES RELIED UPON BY EXAMINER

Appellants wish to initially point out that it is not clear from the final Office Action whether, or to what extent, the Examiner is relying on Bieganski and Robinson in rejecting the claims. In this regard, the Examiner's analysis of the claims (beginning at page 7) refers only to the newly-cited Anderson reference, and does not set forth any basis for combining Anderson with the other two references. Although the Examiner is now relying on a new basis for rejection, Appellants assume the Examiner is relying on Bieganski and Robinson to the same extent they were relied on in the non-final Office Action mailed on January 16, 2003.

Anderson, Bieganski and Robinson (collectively "the applied references") fail to collectively or individually disclose many of the inventive features that are relevant to this appeal. For example, none of the applied references suggests notifying users of an electronic catalog of relationships between specific items and communities, as set forth in many of the claims. In this regard, these references completely fail to recognize that users can make more informed item selection decisions if they are exposed to such information.

The applied references also fail to suggest a process for identifying those items that distinguish the particular community from a general user population (e.g., as the result of the items being much more popular within the community than within the general user population), as set forth in some of the claims. The applied references also fail to suggest using either email addresses or shipping addresses of users to identify specific subsets or communities of users, as set forth in certain claims.

Each reference is discussed in further detail below.

1. Anderson (U.S. Patent No. 5,974,396)

Anderson discloses a system for analyzing consumer purchasing history information to assist a physical retailer in providing targeted advertisements to its customers. The retailer collects demographic information by having its customers fill out membership applications, and enters this data into the system. These customers are also provided with membership cards that allow their purchases to be tracked. See col. 8, lines 21-43. Demographic data supplied by the

Appl. No. : 09/377,447
Filed : August 19, 1999

customers is used to define consumer clusters; for example, a consumer cluster may be defined that consists of all consumers having an income of \$0-25,000, a household size of 1-3, and a head of household age of 20-35 years. See col. 12, lines 10-25. The customers' purchase histories are analyzed in terms of these consumer clusters to identify relationships between customers and products. The results of the analysis are used to print out targeted advertisements or promotions to provide to specific customers. See col. 6, lines 24-48.

The reference does not appear to involve an electronic catalog.

Anderson does not disclose or suggest providing the results of the analysis to the customers to assist such customers in selecting items. Rather, Anderson's sole purpose for analyzing the purchase histories is apparently to provide the retailer with information for marketing purposes. Thus, for example, there is no suggestion within Anderson to notify users of a relationship that has been detected between a particular item and a particular user community.

In addition, Anderson's analysis does not appear to involve the identification of items that distinguish a particular community from a general user population. To the contrary, Anderson's approach appears to merely involve detecting those items that are the most popular within specific groups or clusters of customers, without the regard to the popularity levels of these items among a more general user population.

2. Bieganski (U.S. Patent No. 6,321,221)

Bieganski discloses a collaborative recommendations process in which users having similar profiles are grouped together into "affinity groups" (also called "neighborhoods") for purposes of generating personalized recommendations. For example, all users that share at least two common items within their respective purchase histories may be grouped together to form a particular affinity group. If a particular item appears frequently within the purchase histories of the members of this affinity group, that item may be recommended to those members who have not yet purchased the item. To increase the likelihood that the recommendations will be useful, items that are generally popular are not recommended.

In contrast to Applicants' system, Bieganski's recommendation process does not notify the target user of a relationship that exists between an item and a particular user community. Rather, the item is merely suggested to the user without an explanation as to why it is being recommended.

Appl. No. : 09/377,447
Filed : August 19, 1999

In this regard, there would be little or no reason in Bieganski's system to identify the particular affinity group to the target user, as these affinity groups are not tied to specific companies, organizations, clubs or geographic regions that would be meaningful to the target user. Yet another deficiency in Bieganski's method is that it is apparently incapable of providing meaningful information to a user who does not yet have an interests profile (e.g., has not yet made any purchases).

3. Robinson (U.S. Patent No. 5,790,426)

Robinson discloses automated collaborative filtering algorithms for recommending items to users. These algorithms seek to identify users that have similar preferences based on the item ratings profiles of the users.

Based on the non-final Office Action dated January 16, 2003, the Examiner appears to rely on Robinson primarily because it briefly discloses the use of a "Chi-Square distribution," which is similar in name to the censored chi-square algorithm disclosed in the present application. See Robinson at col. 28, lines 49-56. However, nothing in Robinson suggests that the Chi-Square distribution disclosed therein is used to identify items that characterize or distinguish a community, as set forth in the present application. Rather, the distribution is apparently used to determine a statistical distance between two users that have rated a common item. See Robinson at col. 6, lines 36-51. The reference therefore appears to lack relevance.

C. DISCUSSION OF THE ISSUES ON APPEAL

For the reasons set forth below, Appellants respectfully submit that the obviousness rejection is improper.

1. The Examiner has not identified a motivation to combine the applied references.

The final Office Action does not set forth any basis for combining the newly-cited Anderson reference with the Bieganski and Robinson references. The Examiner therefore has not set forth a prima facie case of obviousness, and the rejection of the claims of Groups 1-6 is improper. See MPEP § 2143.

2. The applied references do not disclose or suggest all of the claim limitations of the broadest claim(s) of each claim group.

In order to establish prima facie obviousness of a claimed invention, all of the claim limitations must be taught or suggested by the prior art. See MPEP § 2143.03. In this case, the applied references fail to satisfy this requirement with respect to the broadest claim or claims of each claim group. Each claim group is discussed separately below.

Group 1

Group 1 consists of Independent Claims 1 and 12 and their respective dependent claims (collectively Claims 1-25). Claims 1 and 12 are reproduced below.

1. A method of assisting users in selecting items from an electronic catalog of items, the catalog accessible to users of an online sales system that provides services for allowing users to purchase items from the catalog, the method comprising:

providing a database which contains information about a plurality of user communities, wherein different communities represent different subsets of users of the sales system;

tracking online purchases of items from the sales system by the users to generate purchase history data, and storing the purchase history data in a computer memory;

processing at least the purchase history data to identify at least one item which, based on actions of both members and non-members of a selected community of said plurality of user communities, has a popularity level that is substantially greater within the selected community than outside the selected community; and

electronically notifying users that the at least one item is popular within the selected community to assist users in selecting items from the electronic catalog.

12. A system for assisting users of an online sales system in selecting items from an electronic catalog of items, the system comprising:

at least one database which contains purchase history data for users of the sales system, and which contains information about a plurality of user communities wherein different communities represent different subsets of users of the sales system; and

a computer process which analyzes at least the purchase history data to identify items that have substantially higher popularity levels within particular communities of the plurality of communities relative to their respective popularity levels among a general user population, and which notifies users of the sales system of the items and associated communities for which such popularity level disparities exist, to thereby assist such users in selecting items from the electronic catalog.

Appl. No. : 09/377,447
Filed : August 19, 1999

With respect to Claim 1, the applied references do not disclose or suggest “electronically notifying users that the at least one item is popular within the selected community to assist users in selecting items from the electronic catalog.” With respect to Claim 12, the applied references fail to disclose or suggest a computer process which notifies users “of the items and associated communities for which such popularity level disparities exist, to thereby assist such users in selecting items from the electronic catalog.” As discussed above, these features of Applicants’ system provide users with valuable information that is helpful to selecting items from the electronic catalog. The Examiner did not acknowledge or address this important deficiency in the applied references.

With respect to Claim 1, the applied references also fail to suggest “processing at least the purchase history data to identify at least one item which, based on actions of both members and non-members of a selected community of said plurality of user communities, has a popularity level that is substantially greater within the selected community than outside the selected community.” With respect to Claim 12, the applied references fail to suggest “a computer process which analyzes at least the purchase history data to identify items that have substantially higher popularity levels within particular communities of the plurality of communities relative to their respective popularity levels among a general user population.” The Examiner did not fully address these limitations of Claims 1 and 12 in the final Office Action.

In connection with these limitations, Bieganski’s recommendation process uses item popularity data to inhibit recommendations of the most popular items, but does not seek to identify those items that have substantially higher popularity levels within a selected community than outside that community. Similarly, Anderson identifies items that are popular among specific clusters of users, but makes no attempt to identify those items that are substantially more popular within a specific cluster than in a general user population. Robinson does not suggest these aspects of the invention by disclosing the use of statistical tools such as a Chi-Square distribution for purposes of generating personalized recommendations, as Robinson’s use of the Chi-Square distribution is for a very different purpose. And none of the cited references suggest disclosing the identity of the community to the user to allow the user to use the community identity as an aid in selecting items.

Appl. No. : 09/377,447
Filed : August 19, 1999

Because the applied references fail to disclose or suggest all of the limitations of Claims 1 and 12, the rejections of the Group 1 claims are improper.

Group 2

Group 2 consists of independent Claim 26 and its dependent claims (collectively Claims 26-31). Claim 26 is directed to an embodiment in which email addresses of users are used to identify a subset of users. The claim reads as follows:

26. A method of assisting users in selecting items from an electronic catalog of items, the catalog accessible to users of an online sales system that provides services for allowing users to purchase items from the catalog, the method comprising the computer-implemented steps of:

identifying a subset of users of the store that have email addresses that satisfy a particular criteria, wherein the subset comprises a plurality of users;

identifying at least one item that characterizes the subset of users, wherein the step of identifying comprises processing purchase history data of the subset of users and of users falling outside said subset; and

electronically notifying users of the online sales system of a relationship between the at least one item and the subset of users, to assist users in selecting items from the electronic catalog.

In connection with this claim and other claims that involve the use of email addresses to identify groups of users, the Examiner asserts that it “is well known that email may be targeted to specific addresses and to particular domain names, perhaps specific to a company or other organization.” The Examiner also asserts that it would therefore have been obvious “to combine Anderson and email to send promotional materials to consumers.” Final Office Action at page 8, last paragraph to page 9, line 8.

Even if these assertions are true (which Appellants do not admit), it does not follow that the method set forth in Claim 26 would have been obvious. In this regard, Claim 26 is not directed to sending promotional materials to consumers. Rather, it is directed to a process that includes “identifying at least one item that characterizes the subset of users [that have email addresses that satisfy a particular criteria],” and “electronically notifying users of the online sales system of a relationship between the at least one item and the subset of users, to assist users in selecting items from the electronic catalog.” The Examiner’s assertions disregard these important claim limitations.

Appl. No. : 09/377,447
Filed : August 19, 1999

Indeed, each of the three subparagraphs of Claim 26 recites a process step that is not disclosed or suggested by the applied references. The rejections of the Group 2 claims are therefore improper.

Group 3

Group 3 consists of independent Claim 51 and corresponding dependent Claims 53-57. Claim 51 is reproduced below, with reference characters for purposes of discussion.

51. A computer-implemented method of assisting users in selecting items from an electronic catalog of items, the method comprising:

- (a) storing email addresses of users of the electronic catalog;
- (b) for each of a plurality of organizations, identifying a respective group of said users who are deemed to be affiliated with the respective organization by virtue of a domain name of the respective organization occurring within an email address of each such user, to thereby identify a plurality of organization-specific groups of users;
- (c) tracking at least one type of user activity that indicates user affinities for particular items in the electronic catalog to generate activity history data for the users;
- (d) for each organization-specific group of users, analyzing the activity history data to generate a list of items within the electronic catalog that distinguish the organization-specific group from a general population of the users, to thereby generate a plurality of organization-specific popular items lists; and
- (e) exposing the organization-specific popular items lists to the users of the electronic catalog in association with the organizations to which such popular items lists correspond, to allow users to identify items that are popular within specific organizations.

Subparagraphs (b), (d) and (e) each include limitations that are not disclosed or suggested by the applied references. With respect to subparagraph (b), for example, none of the applied references involves the use of user email addresses to identify a plurality of organization-specific groups of users, as claimed. Even if Anderson were combined with well known email methods as suggested by the Examiner, the combination would not involve this portion of the claimed process.

With respect to subparagraph (d), neither the targeted promotion methods of Anderson nor the recommendations processes of Bieganski and Robinson involve the analysis of user activity data "to generate a list of items within the electronic catalog that distinguish the organization-specific group from a general population of the users."

Appl. No. : 09/377,447
Filed : August 19, 1999

With respect to subparagraph (e), as discussed above, none of the applied references involves notifying users of items that are popular within specific organizations. Information about the items that are popular within specific organizations can be very helpful to users of the electronic catalog—particularly where these items distinguish members of the organization from a general user population, as claimed.

Because the applied references do not disclose all of the limitations of Claim 51, the rejections of the Group 3 claims are improper.

Group 4

Group 4 consists of dependent Claim 52. Claim 52 depends from Claim 51, and is therefore patentable over the applied references for the reasons set forth above for Group 3.

In addition, Claim 52 recites the additional feature of “providing to a target user a personalized web page that displays organization-specific popular items lists of organizations with which the target user is affiliated.” A preferred embodiment of this feature is illustrated in Figure 2 of the present application. None of the applied references discloses or suggests this feature. If the Examiner disagrees, he is respectfully requested to specifically identify, within his Reply Brief, the reference and the portion thereof he is relying on.

Because the applied references do not disclose or suggest all of the limitations of Claim 52, the rejection of Claim 52 is improper.

Group 5

Group 5 consists of dependent Claim 53, which depends from Claim 51. The rejection of Claim 53 is improper for the reasons set forth above for Group 3.

In addition, Claim 53 recites the additional feature of substantially excluding from consideration activity history data falling outside a selected look-back period. This feature allows the organization-specific popular items lists to be generated based primarily on a most recent set of activity data, so that the lists tend to reflect the current interests of users. As none of the applied references discloses or suggests this feature, the limitations added by Claim 53 provide an additional basis for patentability.

Group 6

Group 6 consists of independent Claim 43 and its dependent claims (collectively Claims 43-50). Claim 43 is directed to an embodiment in which user shipping addresses are used to

Appl. No. : 09/377,447
Filed : August 19, 1999

identify the subset of users. The claim reads as follows, with reference characters added for purposes of discussion:

43. A method of assisting users in selecting items from an electronic catalog of items, comprising:

(a) identifying, among a population of users of the electronic catalog, a subset of users that, based on user shipping addresses, reside in a common geographic region;

(b) monitoring actions of the subset of users, and of users falling outside the subset, to generate user activity history data reflective of user actions performed with respect to items in the electronic catalog;

(c) processing the user activity history data to identify an item whose popularity level among the subset of users is substantially higher than a popularity level of the item among the population of users; and

(d) notifying at least one user that a relationship exists between the item and the geographic region to assist the at least one user in selecting items from the electronic catalog.

Anderson does not suggest the limitations of subparagraphs (c) and (d) of this claim. Specifically, as discussed above, the transaction history analysis of Anderson does not involve the identification of “an item whose popularity level among the subset of users is substantially higher than a popularity level of the item among the population of users,” as recited in subparagraph (c). In addition, nothing in Anderson suggests the notification function set forth in subparagraph (d).

Bieganski and Robinson do not overcome these deficiencies in Anderson.

Because the applied references do not disclose or suggest all of the limitations of Claim 43, the rejections of the Group 6 claims are improper.

3. The applied references do not suggest the desirability of the claimed invention.

As set forth in MPEP 2143.01, in order to establish obviousness based on a combination of references, the prior art must suggest the desirability of the claimed combination. “The mere fact that references can be combined or modified does not render the resultant combination obvious unless the prior art also suggests the desirability of the combination.” MPEP 2143.01 (emphasis original).

Appl. No. : 09/377,447
Filed : August 19, 1999

In the present case, all of the independent claims include language related to the function of notifying users of an electronic catalog of a relationship between a catalog item and a user group or community. As discussed above, this aspect of the invention is desirable as it assists users in making more informed item selection decisions (e.g., by allowing users to identify items that have gained acceptance within specific companies, organizations, geographic locations, etc.). The applied references do not suggest this desirability, but rather focus on identifying items to recommend or promote to users.

Because the applied references fail to suggest the desirability of the claimed invention, the rejections of the claims of Groups 1-6 are improper.


IX. CONCLUSION

For the reasons set forth above, Appellants submit that the rejections of the claims of Groups 1-6 are improper, and request that these rejections be reversed.

Respectfully submitted,

KNOBBE, MARTENS, OLSON & BEAR, LLP

Dated: 11-6-03

By: 
Ronald J. Schoenbaum
Reg. No. 38,297
Attorney of Record
2040 Main Street, 14th Floor
Irvine, CA, 92614
949-721-2950

APPENDIX A – PENDING CLAIMS

1. A method of assisting users in selecting items from an electronic catalog of items, the catalog accessible to users of an online sales system that provides services for allowing users to purchase items from the catalog, the method comprising:

providing a database which contains information about a plurality of user communities, wherein different communities represent different subsets of users of the sales system;

tracking online purchases of items from the sales system by the users to generate purchase history data, and storing the purchase history data in a computer memory;

processing at least the purchase history data to identify at least one item which, based on actions of both members and non-members of a selected community of said plurality of user communities, has a popularity level that is substantially greater within the selected community than outside the selected community; and

electronically notifying users that the at least one item is popular within the selected community to assist users in selecting items from the electronic catalog.

2. The method of Claim 1, wherein electronically notifying users comprises generating a Web page which includes a community-based most popular items list.

3. The method of Claim 2, wherein the most popular items list is a bestsellers list.

4. The method of Claim 1, wherein electronically notifying users comprises automatically generating and sending an email message to members of the selected community.

5. The method of Claim 4, wherein the email message contains contact information of at least one member of the community that has purchased an item described in the email message.

6. The method of Claim 1, wherein processing the purchase history data to identify at least one item comprises identifying a set of characterizing purchases for the community.

7. The method of Claim 1, wherein the community is an implicit membership community.

8. The method of Claim 7, wherein the community is based on email addresses of users.

9. The method of Claim 1, wherein the community is an explicit membership community.

10. The method of Claim 1, wherein the community is derived from an electronic address book of a user.

11. The method of Claim 1, wherein the community is a composite community which comprises multiple other communities of the database.

12. A system for assisting users of an online sales system in selecting items from an electronic catalog of items, the system comprising:

at least one database which contains purchase history data for users of the sales system, and which contains information about a plurality of user communities wherein different communities represent different subsets of users of the sales system; and

a computer process which analyzes at least the purchase history data to identify items that have substantially higher popularity levels within particular communities of the plurality of communities relative to their respective popularity levels among a general user population, and which notifies users of the sales system of the items and associated communities for which such popularity level disparities exist, to thereby assist such users in selecting items from the electronic catalog.

13. The system of Claim 12, wherein the process comprises a first process which generates a table which contains bestselling items lists for at least some of the communities, and a second process which selects items from the table to display to users.

14. The system of Claim 12, further comprising a user interface which allows users to select and join at least some of the user communities.

15. The system of Claim 12, further comprising a user interface which allows a user to define a composite community that includes multiple communities of the database, and to initiate generation of a popular items list for the composite community.

16. The system of Claim 12, wherein at least some of the communities are implicit membership communities.

17. The system of Claim 12, wherein at least some of the communities are based on email addresses of users.

18. The system of Claim 12, wherein at least some of the communities are based on electronic address books of the users.

19. The system of Claim 12, wherein the process generates and displays community bestsellers lists for at least some of the communities.

20. The system of Claim 12, wherein the process compares a popularity of an item within a community to a popularity of the item among non-members of the community.

21. The system of Claim 12, wherein the process sends to the users notification emails that include descriptions of the items that are popular within particular communities.

22. The system of Claim 21, wherein at least some of the notification emails include contact information of users that have purchased items described therein.

23. The system of Claim 21, wherein at least some of the notification emails specify a level of acceptance an item has attained within a particular community.

24. The system of Claim 12, wherein the process identifies items that are have substantially higher popularity levels within particular communities by at least identifying a set of items purchased by members of a community that distinguishes the community from a general user population.

25. The system of Claim 24, wherein the process uses a censored chi-square algorithm to identify the set of items.

26. A method of assisting users in selecting items from an electronic catalog of items, the catalog accessible to users of an online sales system that provides services for allowing users to purchase items from the catalog, the method comprising the computer-implemented steps of:

identifying a subset of users of the store that have email addresses that satisfy a particular criteria, wherein the subset comprises a plurality of users;

identifying at least one item that characterizes the subset of users, wherein the step of identifying comprises processing purchase history data of the subset of users and of users falling outside said subset; and

electronically notifying users of the online sales system of a relationship between the at least one item and the subset of users, to assist users in selecting items from the electronic catalog.

Appl. No. : 09/377,447
Filed : August 19, 1999

27. The method of Claim 26, wherein identifying a subset of users comprises identifying all users having a selected domain name within their respective email addresses.

28. The method of Claim 27, wherein the selected domain name is a domain name of a selected company.

29. The method of Claim 26, wherein identifying a subset of users comprises identifying all users that have one of a selected group of domain names within their respective email addresses.

30. The method of Claim 26, wherein electronically notifying comprises generating a Web page which includes a list of items that characterize the subset of users.

31. The method of Claim 26, wherein electronically notifying comprises sending email notification messages to at least some of the users of the subset.

32-42 (not involved in this appeal)

43. A method of assisting users in selecting items from an electronic catalog of items, comprising:

identifying, among a population of users of the electronic catalog, a subset of users that, based on user shipping addresses, reside in a common geographic region;

monitoring actions of the subset of users, and of users falling outside the subset, to generate user activity history data reflective of user actions performed with respect to items in the electronic catalog;

processing the user activity history data to identify an item whose popularity level among the subset of users is substantially higher than a popularity level of the item among the population of users; and

notifying at least one user that a relationship exists between the item and the geographic region to assist the at least one user in selecting items from the electronic catalog.

44. The method as in Claim 43, wherein processing the user activity history data to identify an item whose popularity level among the subset of users is substantially higher comprises at least one of (a) comparing a velocity of the item within the subset to a velocity of the item within a general user population; and (b) comparing an acceleration of the item within the subset to an acceleration of the item within a general user population.

Appl. No. : **09/377,447**
Filed : **August 19, 1999**

45. The method as in Claim 43, wherein processing the user activity history data to identify an item whose popularity level among the subset of users is substantially higher comprises applying a censored chi-square type algorithm to the user activity history data.

46. The method as in Claim 43, wherein notifying at least one user comprises presenting to the user a list of items that are popular in the geographic region relative to popularity levels outside the region.

47. The method as in Claim 43, wherein the user actions performed with respect to items in the electronic catalog comprise item purchases.

48. The method as in Claim 43, wherein the user actions performed with respect to items in the electronic catalog comprise item viewing events.

49. The method as in Claim 43, wherein the geographic region corresponds to a particular city.

50. The method as in Claim 43, wherein processing the user activity history data comprises applying a time window to the user activity history data such that a popularity level of the item is measured substantially from a most recent set of the user activity history data.

51. A computer-implemented method of assisting users in selecting items from an electronic catalog of items, the method comprising:

storing email addresses of users of the electronic catalog;

for each of a plurality of organizations, identifying a respective group of said users who are deemed to be affiliated with the respective organization by virtue of a domain name of the respective organization occurring within an email address of each such user, to thereby identify a plurality of organization-specific groups of users;

tracking at least one type of user activity that indicates user affinities for particular items in the electronic catalog to generate activity history data for the users;

for each organization-specific group of users, analyzing the activity history data to generate a list of items within the electronic catalog that distinguish the organization-specific group from a general population of the users, to thereby generate a plurality of organization-specific popular items lists; and

Appl. No. : 09/377,447
Filed : August 19, 1999

exposing the organization-specific popular items lists to the users of the electronic catalog in association with the organizations to which such popular items lists correspond, to allow users to identify items that are popular within specific organizations.

52. The method as in Claim 51, wherein exposing the organization-specific popular items lists to the users comprises providing to a target user a personalized web page that displays organization-specific popular items lists of organizations with which the target user is affiliated.

53. The method as in Claim 51, wherein analyzing the activity history data comprises substantially excluding from consideration activity history data falling outside a selected look-back period.

54. The method as in Claim 51, wherein analyzing the activity history data comprises applying a censored chi-square type algorithm to the activity history data.

55. (Canceled)

56. The method as in Claim 51, wherein at least some of the organizations are companies.

57. A computer system that embodies the method of Claim 51.

Visualizing Amazon.com's Purchase Circles

David DeMumbrum, Jeannie Lee, William Stark, Bessie Yang

College of Computing

Georgia Institute of Technology

Atlanta, GA 30332 USA

{ddemumbr, jeannie, wstark, bessie}@cc.gatech.edu

ABSTRACT

Amazon is a well-known online company that sells products such as books and music. It also tracks the purchasing patterns of a variety of groups including private corporations, government organizations, and geographic areas. Amazon defines each of these groups as a "purchase circle." For each purchase circle, Amazon lists the best-selling items in the Books, Music, Video, DVDs, and Electronics product categories. Our objective is to create a dynamic visualization of Amazon's purchase circles that focuses on looking at the Top 10 music titles and genres that are popular in selected U.S. cities. We present a visualization known as CityPrints, a dynamic query-based tool for producing color-coded visual representations of purchase circles data. CityPrints allows users to quickly compare popular titles in different U.S. cities, identify which music genres are popular in a given city, and rank cities according to how popular a given music genre is in that city.

Keywords

information visualization, purchasing patterns, visualization, dynamic queries

INTRODUCTION

Amazon is a well-known online company (<http://www.amazon.com>) that sells a wide variety of products such as books, movies, and music online. One of its more interesting characteristics is a purchase circles feature which tracks the purchasing patterns of a variety of groups including private corporations, non-profit and professional organizations, colleges and universities, government organizations, and geographic locations. Amazon defines each of these groups as a "purchase circle."

For each purchase circle, Amazon lists the best-selling items in the Books, Music, Video, DVDs, and Electronics product categories, but it does not rely solely on overall sales when compiling its lists. Instead, Amazon creates its purchase circle lists by aggregating data on items that have

been shipped to a particular location or ordered from a particular domain name. It then modifies this information to create lists of items that are more popular for that group in comparison with the general population. The update frequency of purchase circles varies from weekly to monthly, depending on the size of the purchase circle [1]. Currently, users who are interested in looking at a specific purchase circle must go to Amazon's Web site and navigate to the desired purchase circle and product type. This results in a list containing information about the Top 10 items that are popular for that purchase circle. For example, in music items, the Top 10 list contains information about the title of a music album, the artist's name, pricing, availability, and average customer rating. With this current system, users have no way of performing side-by-side comparisons of multiple purchase circles or quickly identifying what types of music are popular in a given purchase circle.

Our dynamic visualization tool, known as CityPrints, is designed to provide these features using an interface that is simple and straightforward. For demonstration purposes, our initial version of the CityPrints tool is focused on the category of music and depicts purchase circle data for a limited number of U.S. cities. Through a list box and an assortment of radio buttons, the user selects a set of cities (numbering between one and five) and a sorting method. In response, the CityPrints tool displays for each selected city a vertical stack of ten horizontal colored bars, one for each of the top ten most popular music selections for that city. Each color corresponds to one of ten unique musical genres. Clicking on any one individual color bar results in the details of the represented selection being displayed in a text window to the left of the set of CityPrints, including the selection's title, artist, genre, year of release, and sales rank.

Our target users are people who are interested in the following kinds of tasks:

- Identify the most popular music title(s) in a specific U.S. city
- Given a specific music title/genre, identify the cities in which it is popular
- Rank the popularity of music genres for selected U.S. cities

- Compare popular titles/genres between multiple U.S. cities

In the sections that follow we will describe our design decisions, related work that contributed to the development of CityPrints, an overview of the CityPrints visualization tool, and finally aspects we would like to focus on during future development of CityPrints.

DESIGN AND IMPLEMENTATION

Our group hypothesized that Amazon's purchase circles data would confirm common presumptions about the popularity of musical genres within certain cities. For example, we expected that Nashville, TN, would feature more country music titles in its Top 10 list than most other U.S. cities. We also hypothesized that some cities would share matched titles in their Top 10 lists. We originally hoped to look more closely at such matches in order to determine if cities containing matched items held additional musical preferences in common. However, as we gathered data from the Amazon Web site, we realized that because very few cities shared matched music titles in their Top 10 lists, there was little opportunity to explore overlapping music tastes by looking only at matched titles. We consequently decided to focus on visualizing comparisons between cities based on musical genre, while providing details on demand for specific albums. To keep our visualization simple, we limited the number of genres to the following ten: Blues/Funk, Classical, Country/Folk, Dance/Electronica, Jazz, Miscellaneous, Rap/R&B, Rock/Pop, Soundtracks, and World.

Displaying Cities

For our dynamic visualization, we wanted to enable our users to compare the music preferences of different cities using as little cognitive effort as possible. Therefore, we chose to capitalize on the human perceptual ability to process certain attributes such as color pre-attentively [7].

In CityPrints, each city is represented by its own CityPrint, a graphical column consisting of 10 colored rows. Each of the rows corresponds to a music title in the city's Top 10 list and is color-coded according to each title's associated genre. For example, the CityPrint in Figure 1 for Chicago contains six red rows as well as one orange, green, pink, and blue row for a total of ten rows. Red rows represent Rock titles, orange represents a Dance/Electronica title, green represents a Miscellaneous title, pink represents a Classical title, and blue represents a Jazz title. This method of encoding provides a simple and quick way for users to see how genres are distributed throughout a city's Top 10 list as well as compare multiple cities to see if they exhibit preferences for particular musical genres.

Chicago




Figure 1. CityPrint for Chicago

One drawback to this technique pertains to its heavy reliance on color as an information-encoding mechanism. In order to make proper use of CityPrints, users must differentiate amongst different colors in the display. CityPrints would therefore be ill suited for use by colorblind individuals and/or users working on a monochrome display. A simple resolution to this problem might involve adding text labels to CityPrint rows in order to specify the genre represented by each row. Users forced to rely upon such labels would miss the pre-attentive processing advantages afforded by color encoding, but they would nevertheless have access to all information within the display. Another potential issue is the fact that while users can quickly discern how many genres are represented in a particular city, they might not be able to identify those specific genres without referring to a legend first. However, we anticipate that this would not be a serious issue, especially for experienced users. A simple feature that could help alleviate this problem would also be to allow users to customize the colors for each music genre according to their own preferences.

User Queries

In order to provide users with continuous updates to specified queries, we based the interactive properties of CityPrints on the dynamic queries pioneered by Shneiderman [6]. However, the limited-case nominal structure of data visualized by CityPrints suggested an alternative to the slider control mechanism classically associated with dynamic query applications. CityPrints instead provides users with a quick means of obtaining custom visualizations based on their use of buttons and lists to select the names of cities and musical genres they wish to see represented. For example, if a user wants to compare music preferences in New York City, NY, and Seattle, WA, she can select those cities within a list, and their corresponding CityPrints will appear on the display area of the visualization, shown in Figure 2.




Figure 2. CityPrints interface with five cities selected

Alternatively, a user interested in determining which cities most strongly prefer a particular type of music can pick a genre (e.g., "Jazz") from the appropriate list and obtain a collection of CityPrints corresponding to the five cities in which the selected genre is most popular. CityPrints displayed as the "Top 5 most popular" for this purpose would be automatically selected by an application-based algorithm that performs calculations to determine which five cities contain the largest overall number of Jazz albums in the CityPrints data set. Only cities with at least one jazz album can be displayed, and if there are fewer than five cities that meet this criterion then only those cities are displayed.

We also considered implementing a variation on this scheme that would choose a Top 5 collection based on what we call a popularity index. In this scheme, the Top 5 CityPrints chosen for display would be based not only on the total number of albums in a particular genre featured within a city's Top 10 list, but also on the overall rank number assigned to each album (a figure Amazon computes based on sales volume). Complexities associated with implementing this idea, combined with a minimal impact on the general interest value of our data set, caused us to bypass this idea for the time being.

While the present CityPrints interface does not quite fit the true definition of dynamic queries due to the need for a user to press a button in order to refresh the display, future versions of our prototype would allow users to see their query results immediately.

Displaying Details

One challenge present throughout the development of CityPrints concerned its ability to deliver detailed information (for example, album title, artist, and so forth) about all of the music titles represented in the display. We

considered using a focus + context method, similar to the Table Lens [5], of expanding a selected row (representing a music title) to show the details. Deselecting a row would cause it to return to a normal-sized colored bar, and multiple rows could be selected at a time. This would avoid requiring users to manage multiple windowed displays. However, the disorienting liabilities associated with this focus + context method could potentially inhibit the user's ability to formulate comparisons of music titles by rank between various cities.

We eventually chose an overview and detail approach where users may obtain "details on demand" for an album by clicking on any row within a particular CityPrint. To facilitate construction of a prototype, we placed the mouse-over details into a tiled window that appears in the upper-left corner of the CityPrints application. Thus, within the current prototype, selected CityPrints comprise the main display, while details associated with individual albums appear inside smaller fixed window off to the side, similar to the tiled multilevel browser described in [4]. This provides a simple solution to the detail problem, but at the same time adds drawbacks of taking up valuable screen space, as well as dividing a detail-seeking user's attention between two separate displays. Another disadvantage when compared to the focus + context method described previously is that the details for only one music title can be displayed at a time.

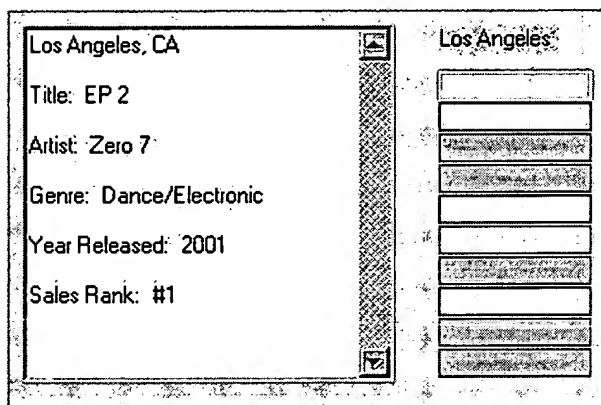


Figure 3. Details for the Los Angeles purchase circle's top music title

Sorting

The default sorting of each CityPrint follows the Amazon-derived ranking of a city's Top 10 music titles by sales volume. Thus, in any CityPrint, the number one selling album on a city's Top 10 list appears by default in the topmost row, while the number 10 selling album appears in the bottom row. One problem is that users might find it

difficult to gather meaningful information using this approach, especially if they are not interested in the relative rankings of music titles within a particular city. For example, if a user wants to determine which city has the most Rock/Pop albums in its Top 10 list, he may have to resort to manually counting each of the red-colored bars which are scattered across each of the CityPrints. To address this, CityPrints' sorting features allow users to modify the ordering of rows within each CityPrint. Users can choose to sort each CityPrint by the default sales ranking, user-specified genre, and overall frequency of each music genre.

Users most interested in particular types of music can utilize the sort feature to modify the ordering of titles according to a specific genre. In this way, the sort feature allows users to automatically cluster together albums within particular genres of interest. For example, if a user chooses to sort by the Country/Folk music genre, then all albums within the Country/Folk genre collect at the top of all visible CityPrints, creating something of an inverted bar chart in Figure 4. In addition to clustering, this technique also allows users to explore particular genres of interest within the context of a genre-based display. For example, the music titles that do not fall under the Country/Folk genre are also clustered together by their corresponding genre, and the genre clusters are ordered alphabetically from top to bottom by the name of the genre (so Classical would appear below Country/Folk which is the selected music genre, but above Dance/Electronica).

In the example shown in Figure 4, users can see that Nashville has the most Country/Folk albums, but they can also quickly see that Atlanta has the most Rock/Pop albums, followed by Dallas, while Nashville has no Rock/Pop albums at all in its Top 10 list. Note that this type of "dual" bar chart occurred because none of the three cities had Soundtracks or World music in their Top 10 lists. However, even when this is not the case, users can still quickly compare the relative widths of each colored genre cluster to determine its popularity for each city without necessarily having to sort by each genre.

Users can also choose to sort selected CityPrints by the overall frequency of each music genre in their respective Top 10 lists. Each CityPrint is sorted with the highest frequency genre clustered at the top, followed by the genre with the next highest number of music titles in the Top 10 list, and so on. This provides users with an alternative method of determining which music genres are most popular for particular cities. For example, in Figure 5, users can quickly see that Rock/Pop is the most popular music genre in Boston with five music titles, followed by

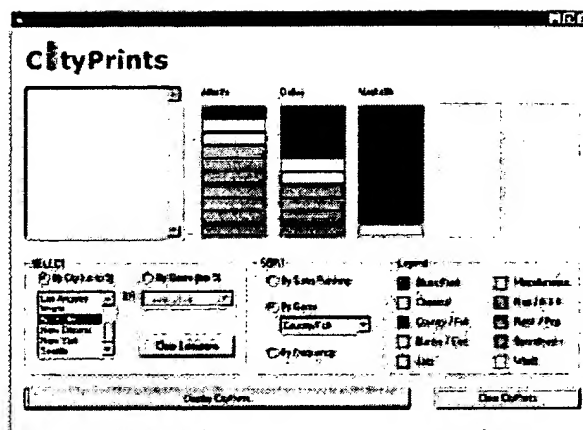


Figure 4. Selected CityPrints sorted by the Country/Folk genre

Classical and Dance/Electronica with two music titles each. In the case of a tie, the genres are placed in alphabetical order from top to bottom.

Unlike the Sort by Genre option, which aids the task of identifying how popular a specific music genre is for selected cities, the Sort by Frequency option aids the task of ranking the popularity of all of the music genres for the selected cities.

By using the sorting features, users can view the relative popularity of a chosen genre across multiple cities while maintaining the ability to view other genres popular in all cities on display. This approach is similar to the work of

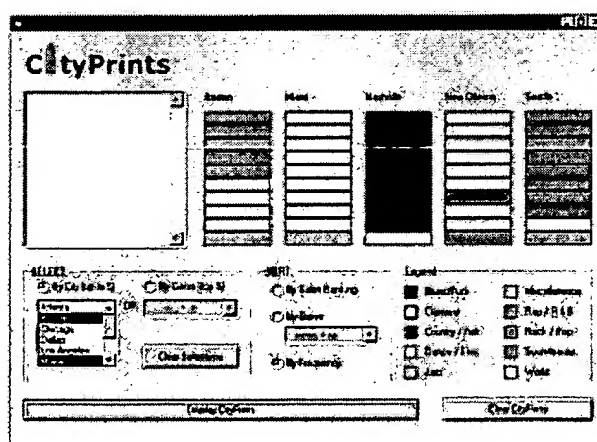


Figure 5. Selected CityPrints sorted by their most popular music genres

Keim et al. [3], whose pixel bar charts also relied heavily upon color as a means of encoding data and used equal-height bars in order to use screen space efficiently. CityPrints differs from pixel bar charts in that it uses graphical rows rather than pixels to represent data items. Since it currently involves only nominal data, CityPrints also does not use the pixel bar chart's equal-height/varied-width mechanism to represent quantitative data.

Implementation

CityPrints was implemented in Visual Basic under various versions of the Windows operating system and currently runs on most versions of Windows. This approach was taken to minimize development time and arrive at a workable prototype rapidly.

The implementation currently has no database backend. All CityPrints and selection details were hard coded for the purpose of demonstrating application capabilities.

For future prototypes, we intend to port the implementation to Java to ensure that the program is platform independent and further employ the use of 2D graphics libraries to reduce image loading and processing overhead during program execution.

FUTURE DIRECTIONS

Although our current implementation of CityPrints does meet the general objective of expressing a city's music preferences in a simple and straightforward manner, we have identified a number of potential improvements.

First and foremost, we would like future versions of CityPrints to include a completed implementation of a mouse-over feature for displaying music title details. Such an improvement would resolve some of the more irksome problems we encountered using overview and detail, in the sense that users could seamlessly obtain information associated with individual music titles without having to consult an extra display off to the side. We envision such a feature being similar to a "gloss" in the Fluid Links system [8], in that it would show additional information without consuming screen space or occluding the primary CityPrints display.

Scalability is another issue we plan to address in future versions which applies to both displaying data and entering queries. In its current form, CityPrints allows users to display up to 5 cities at a time, with each city arranged in alphabetical order from left to right. We would like to allow users the ability to select and view more cities simultaneously, as well as enable them to dynamically change the position of a CityPrint instead of locking it into a predetermined slot. Users could then cluster CityPrints in a flexible manner without being forced into a grid-based,

alphabetically arranged format. The size of each CityPrint could also be dynamically changed to accommodate the number of displayed cities, although this would require careful consideration of how large each CityPrint must be to be useful.

In terms of the scalability of entering user queries, users currently select cities using a dropdown box which, while acceptable for our current prototype that includes only 10 cities, would complicate selection tremendously if we significantly increased the number of available locations. One possible solution might involve using a map to display a larger number of cities. This approach would provide users with a natural and intuitive means of city selection (provided they were familiar with the geography of the United States), and also contribute the added variable of geographic location to possibly enrich the overall display of purchase circle information. For example, with a map, users could better consider the possible impact of geographic location on a city's musical preferences. A map-based representation would complement our slated future implementation of dynamic query mechanisms. For example, it could allow users to request a display of CityPrints that contain at least two albums belonging to the Blues/Funk genre.

Our consideration of such ideas was influenced by the work of Dang, North, and Shneiderman [2], whose Dynamaps made use of dynamic queries through brushing techniques applied to choropleth maps. Our design would take a similar approach in that users can specify a query that causes the resulting cities to light up in the map display and the corresponding CityPrints to appear in an adjacent display. A more advanced version of this might allow users to select cities by highlighting a rectangular region of the map (for example, selecting only cities in the Southeast United States) with the mouse cursor and then entering queries that apply only to the selected cities. CityPrints would differ from Dynamaps in that we would concentrate only on the display of cities for selection purposes rather than using a choropleth map.

A drawback associated with using a map selection scheme concerns the nontrivial difficulty of differentiating amongst (and potentially labeling) cities in close geographic proximity to one another (for example, Chicago, IL, and Gary, IN). A poorly implemented map-based selection mechanism could require the user to execute some extremely tricky mouse movements! A zooming interface could ease this problem somewhat; however, we would carefully need to weigh the positive and negative aspects of such an elaborate user interface work in order to avoid confusing users and cluttering the CityPrints application environment.

In addition to an enhanced city selection mechanism, we would like to further enhance the existing CityPrints prototype to display time series information associated with each city. We found interesting the prospect of allowing users to view a city's current CityPrint and compare it with CityPrints of the past (for example, 6 months ago, 1 year ago, 2 years ago, and so on). Such a feature would allow users to view how location-based music preferences may have shifted over time.

CONCLUSIONS

We have designed and implemented a dynamic visualization tool named CityPrints that allows users to view and compare music preferences in selected U.S. cities, as reported in Amazon's purchase circles. Our objective was to organize and display a city's music preferences using a simple and straightforward interface. We feel that we have achieved this objective with our current implementation, which allows users to select cities by name and popular music genres as well as sort each CityPrint according to sales ranking, user-specified genre, and overall frequency of each genre in a CityPrint. However, there are still a number of potential improvements and added features that can improve the functionality and usefulness of CityPrints. These include the use of semi-transparent windows to display information about selected music titles, enhanced scalability, geographic representation of cities, selection of cities in a manner more closely in-line with the dynamic queries paradigm, and the incorporation of time series data.

REFERENCES

1. Amazon.com Purchase Circles FAQ. Available at <http://www.amazon.com/exec/obidos/tg/stores/browse/-/help/468604/102-9651313-6828964>.
2. Dang, G., North, C., and Shneiderman, B. Dynamic queries and brushing on choropleth maps, in *Proceedings of the Fifth International Conference on Information Visualization 2001* (London UK, July 2001), 757-764.
3. Keim, D., et al. Pixel bar charts: A new technique for visualizing large multi-attribute data sets without aggregation. *IEEE Symposium on Information Visualization 2001* (October 2001), 113-120.
4. Plaisant, C., Carr, D., and Shneiderman, B. Image-browser taxonomy and guidelines for designers. *IEEE Software* 12, 2 (March 1995), 21-32.
5. Rao, R., and Card, S.K. The Table Lens: Merging graphical and symbolic representations in an interactive focus + context visualization for tabular information, in *Proceedings of the ACM SIGCHI '94 Conference on Human Factors in Computing Systems* (Boston MA, April 1992), ACM Press, 318-322.
6. Shneiderman, B. Dynamic queries for visual information seeking. *IEEE Software* 11, 6 (1994), 70-77.
7. Woods, D. The alarm problem and directed attention in dynamic fault management. *Ergonomics* 38, 11 (1995), 2371-2394.
8. Zellweger, P., Chang, B., and Mackinlay, J. Fluid links for informed and incremental link transitions, in *Proceedings of Hypertext '98* (Pittsburgh PA, June 1998), 50-57.